An Adaptively Constructed Algebraic Multigrid Preconditioner for Irreducible Markov Chains ☆

James Brannick[b], Karsten Kahl[a], Sonja Sokolovic[a]

[a]*Department of Mathematics, Bergische Universität Wuppertal, 42097 Germany*
[b]*Department of Mathematics, Pennsylvania State University, University Park, PA 16802, USA*

arXiv:1402.4005v1 [math.NA] 17 Feb 2014

Abstract

The computation of stationary distributions of Markov chains is an important task in the simulation of stochastic models. The linear systems arising in such applications involve non-symmetric M-matrices, making algebraic multigrid methods a natural choice for solving these systems. In this paper we investigate extensions and improvements of the bootstrap algebraic multigrid framework for solving these systems. This is achieved by reworking the bootstrap setup process to use singular vectors instead of eigenvectors in constructing interpolation and restriction. We formulate a result concerning the convergence speed of GMRES for singular systems and experimentally justify why rapid convergence of the proposed method can be expected. We demonstrate its fast convergence and the favorable scaling behaviour for various test problems.

Keywords: Markov chains, algebraic multigrid, bootstrap AMG, adaptive methods, least squares interpolation, singular vectors, preconditioned GMRES iteration, nonsymmetric systems

1. Introduction

Given the nonsymmetric transition matrix A of a discrete Markov chain, the task is to compute the steady state vector $x \neq 0$ of the Markov chain such that

$$Ax = x, \tag{1}$$

corresponding to the eigenvector of A with eigenvalue 1. If the underlying Markov chain defining the matrix A is irreducible, then the existence and uniqueness (up to a scalar factor) of x can be guaranteed by the Perron-Frobenius theorem (cf. [1]). To compute the steady state vector x, we reformulate (1) as a (singular) linear system of equations

$$Bx = 0, \text{ where } B = I - A \tag{2}$$

as has been proposed in, e.g., [2, 3]. We then adaptively construct an algebraic multigrid hierarchy for the matrix B by means of the *bootstrap algebraic multigrid framework* (BAMG) [2, 4, 5, 6]. In contrast to the BAMG approach considered in [2], which aimed at approximating eigenvectors of A in the bootstrap setup, we now take left and right singular vectors of A as test vectors for the construction of the transfer operators of the multigrid hierarchy. This approach has been developed for nonsymmetric M matrices coming from discretization of convection diffusion problems [7] and has been shown to be effective in practice for some problem classes. As we show here, such a strategy is in addition well suited for transition matrices coming from Markov chains.

In a second phase of the proposed method, the adaptive process is exited and the existing multigrid hierarchy is used to formulate a preconditioner for a GMRES solver applied to the residual equation, computed from the approximation x obtained in the first (setup) phase. Since the system matrix B is singular, the convergence of the

☆The work of the first author was partially supported by the National Science Foundation grant DMS-1320608 and by a U.S. Department of Energy subcontract B605152.

Email address: sokolovic@math.uni-wuppertal.de, +49 202 439 3776 (Sonja Sokolovic)

preconditioned GMRES iteration cannot be guaranteed in general, but we find that in practice the preconditioned method convergences for a wide variety of test problems.

While the algorithm we develop can be viewed as a general approach for solving nonsymmetric problems, we focus here on nonsymmetric and singular M-matrix systems arising in Markov chain applications. Important applications in which one needs the steady state solution of a Markov chain include statistical mechanics [8], queuing theory [9], analysis of telecommunication networks [10], information retrieval [11] and web ranking (e.g., Google PageRank) [12, 13]. Many different methods for computing the steady state distribution have been proposed in the literature. These include the so called aggregation/disaggregation methods. These can be related to algebraic multigrid methods in the sense that they use a series of increasingly smaller Markov chains which are generated by aggregating groups of states into a single state. Most aggregation/disaggregation methods are two-level methods (cf. [14, 15, 16]) as they are based on the original approach in [17, 18], and it took approximately twenty years until multi-level aggregation/disaggregation methods were proposed [19, 20]. Other, more recent, research on the design of multilevel approaches for Markov chains include the development of methods related to the *smoothed aggregation multigrid* framework [12, 3, 21], a Schur complement based multilevel approach used as a preconditioner for the GMRES iteration [22], methods based on iterant recombination with minimization in the ℓ_1- or ℓ_2-norm [23, 24], and a bootstrap method [2] which serves as the starting point for the new method presented in this paper.

The remainder of this paper is organized as follows: In section 2 and 3, we provide an overview of basic material concerning Markov chains and multigrid methods, respectively. In section 4 we give a review of the idea and the various components of the bootstrap algebraic multigrid framework and provide details concerning the use of singular vectors as test vectors. Some special adjustments of the framework, making it suitable for dealing with Markov chains, are described in section 5. In section 6, details on how to use the BAMG hierarchy as a preconditioner for a GMRES iteration are given, and the speed of convergence of the resulting method is investigated. Several numerical examples illustrating the performance of the resulting method are given in section 7. In section 8, concluding remarks and an outlook on topics for future research are given.

2. Markov chains

A discrete finite Markov chain with states $\{1, \ldots, n\}$ and transition probabilities $p_{ij} \geq 0$ between these states can be identified with a matrix $A \in \mathbb{R}^{n\times n}$, the transition matrix, by setting $a_{ij} = p_{ji}$. By this definition the transition matrix is column stochastic, i.e., $a_{ij} \geq 0$ for all i, j and the column sums $\sum_{i=1}^{n} a_{ij}$ of A are one for all columns j. In addition, we will assume that the matrix A is *irreducible*, i.e., there exists no permutation matrix Π such that

$$\Pi^T A \Pi = \begin{bmatrix} A_{11} & 0 \\ A_{21} & A_{22} \end{bmatrix}.$$

A geometric interpretation of irreducibility is that the directed graph induced by A is strongly connected. For an irreducible transition matrix A the Perron-Frobenius theorem (cf. [1]) implies that the system

$$Ax = x \tag{3}$$

always has a unique solution x (up to scaling) with all its entries strictly positive.

To determine the steady state vector that solves (3) by an algebraic multigrid method, we first reformulate the eigenproblem as a linear system

$$Bx = 0, \text{ where } B = I - A. \tag{4}$$

Since 1 is a simple eigenvalue of A, the matrix B is singular and $\operatorname{rank}(B) = n - 1$.

3. Algebraic multigrid

Before describing the bootstrap algebraic multigrid framework in section 4, we give a very brief review of the basic concepts and individual components of a multigrid algorithm. The first ingredient of any multigrid method is the *smoothing scheme*, which is often a stationary iterative method based on a splitting of the matrix B, e.g., weighted

Richardson, weighted Jacobi, Gauss-Seidel [25] or Kaczmarz relaxation [26]. Any stationary smoothing iteration is based on a matrix splitting $B = F - G$ with F non-singular, and the smoothing iteration for a system $Bx = b$ reads

$$x^{(k+1)} = F^{-1}Gx^{(k)} + F^{-1}b =: \mathcal{S}(x^{(k)}, b), \quad k = 0, 1, \ldots. \tag{5}$$

Its error propagation is given by

$$e^{(k+1)} = S\, e^{(k)} \text{ with } S = I - F^{-1}B, \tag{6}$$

where $e^{(k)} = x - x^{(k)}$ is the error of the k-th iterate $x^{(k)}$. For weighted Jacobi, e.g., we have $F = \omega D$, with D the diagonal of B. The general idea of multigrid methods is based on the observation that smoothing acts locally (it combines values of variables which are neighbors in the graph of the matrix) and it therefore tends to eliminate certain error components, the "local" errors, very fast, whereas other, "global", components remain almost unchanged. This results in slow overall convergence after an initial phase where the norm of the error is reduced substantially. For these local iterative smoothing schemes the ansatz is made that the error components e which cause slow convergence satisfy the inequality

$$\|Be\| \ll \|e\|, \tag{7}$$

and they are referred to as *algebraically smooth* error. Typically, in a multigrid method, a few smoothing iterations are first applied to substantially reduce the algebraically non-smooth error components. The residual is then restricted to a subspace of smaller dimension, the *coarse grid*, using a restriction operator Q, where the remaining algebraically smooth error can be treated more efficiently. After computing an approximate representation of this error in the subspace from the restricted residual, using a suitably defined *coarse grid operator* B_c, this error representation is interpolated back to the original fine space using an interpolation operator P. It is then added as the *coarse grid correction* to the current approximation there, and the thus corrected approximation is smoothed again. If the linear system on the coarse grid is solved exactly, the error propagation due to the coarse grid correction process is given by

$$e^{(k+1)} = e^{(k)} - PB_c^{\dagger}Qr^{(k)}, \tag{8}$$

where $r^{(k)} = Be^{(k)} = -Bx^{(k)}$ is the residual of (4) and $B_c^{\dagger}$ denotes the Moore-Penrose pseudo inverse of B_c, see, e.g., [27]. Including ν_1 pre- and ν_2 postsmoothing iterations, the error propagation matrix of the entire two-grid method is therefore

$$E = (I - F^{-1}B)^{\nu_1}(I - PB_c^{\dagger}QB)(I - F^{-1}B)^{\nu_2}. \tag{9}$$

Recursive application of this approach (instead of solving exactly on the coarse grid) gives rise to a multigrid (instead of a two-grid) method.

In this paper, for the sake of simplicity, we only consider V-cycles in the solve phase, i.e., we solve the coarse-level equation only once on each level within each multigrid iteration, but the method we propose can easily be generalized to incorporate other cycling strategies, cf. [28, 29]. In the following, we always denote the restriction operator by Q, the interpolation operator by P, and we assume that the coarse grid subspace is given by a subset of the variables of the original fine space. We denote the sets of coarse and fine (and not coarse) variables by C and $\mathcal{F}$, respectively, with $C \cup \mathcal{F} = \{1, \ldots, n\}$ and $C \cap \mathcal{F} = \emptyset$. This choice of coarse grid variables is often referred to as a $C/\mathcal{F}$-splitting (cf. [30]). As coarse grid operator we choose the Petrov-Galerkin operator $B_c = QBP$. Whenever possible a two-grid notation is used for notational simplicity, with n_c standing for the number of coarse variables. In cases where we need to consider all levels of the hierarchy we number them from 1 to L, where 1 denotes the finest grid and L the coarsest.

4. Bootstrap algebraic multigrid (BAMG)

In this section, we review the components of the basic BAMG framework and discuss the incorporation of singular vectors as test vectors in the adaptive process. The BAMG framework [4, 5] is a fully adaptive algebraic multigrid method for the solution of general linear systems. It is based on three main components:

1. compatible relaxation to determine the set of coarse variables C,
2. least squares based computation of the transfer operators, and
3. adaptive construction of appropriate test vectors used to build the transfer operators.

This paper focuses on the second and third components of the BAMG process. We refer to [31, 32, 33] for details on compatible relaxation and to [2] for its application to Markov chain problems.

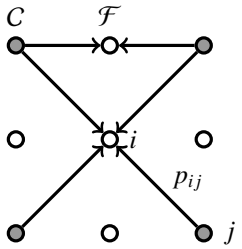

Figure 1: Interpolation naming convention: The $\mathcal{F}$-variable i is interpolated from the four gray C-variables which form the set $\mathcal{J}_i$.

4.1. Least squares based interpolation

Given a set of coarse variables, e.g., computed by compatible relaxation, or resulting from geometric coarsening, the next task in the AMG setup phase is to compute a restriction operator, Q, and an interpolation operator, P, to transfer residuals from fine to coarse grids and to transfer corrections from coarse to fine grids. A careful choice of these operators is crucial to the design of an efficient multigrid method.

In classical multigrid theory (cf. [30]) the *(strong) approximation property* is used to guide the construction of intergrid transfer operators that yield an efficient method. The approximation property can be formulated as the condition

$$\min_{v_c} \|v - Pv_c\|_B^2 \leq \frac{K}{\|B\|_2} \langle Bv, Bv \rangle \text{ for a small constant } K > 0, \tag{10}$$

for each v on the fine grid when B is assumed to be positive definite. Its essence is that a vector v needs to be approximated in the range of P with an accuracy that is inversely proportional to $\|Bv\|$. This is motivated by the fact that vectors v with $\|Bv\| \ll \|v\|$ (i.e., algebraically smooth vectors) dominate the error after applying the smoothing iteration and must therefore be reduced via the coarse grid correction. If no a priori information on algebraically smooth vectors is known, a common idea for adaptively constructing suitable transfer operators is to gather information on algebraically smooth vectors in terms of *test vectors* and then determine P so that these test vectors lie in or near the range of P. In the context of BAMG, this is done by determining an interpolation operator that approximates a set of given vectors in a least squares sense. Specifically, given a set of test vectors $\mathcal{V} = \{v^{(k)}, k = 1, \dots r\}$ such that $\|v^{(k)}\|_2 = 1$, a weighted least squares problem of the form

$$\min \mathcal{L}(P_i) = \sum_{k=1}^{r} \omega_k \Big(v^{(k)} - \sum_{j \in \mathcal{J}_i} (P_i)_j (R v^{(k)})_j \Big)^2 \tag{11}$$

is solved for each row $P_i = (p_{ij})_{j \in \mathcal{J}_i}$ of the prolongation operator P belonging to a fine-level variable $i \in \mathcal{F}$. In (11) R denotes the canonical injection which maps every vector onto its coarse grid components and $\mathcal{J}_i$ denotes the index set of (coarse) variables from which the i-th variable interpolates, i.e., the non-zero pattern of row P_i, cf. Figure 1. The interpolation acts as the identity on the coarse grid variables C such that, when ordering the $\mathcal{F}$ variables first, P has the structure

$$P = \begin{bmatrix} P_{fc} \\ I \end{bmatrix}. \tag{12}$$

The sets $\mathcal{J}_i$ can be determined using the geometry of the problem or can be chosen from a local neighborhood of i in the graph corresponding to B using a greedy algorithm, see [2, 4]. When an algebraic approach is used, a threshold c for the number of interpolation variables has to be set to ensure the sparsity of the interpolation operator and, thus, to limit the growth of the operator complexity

$$o_c = \frac{1}{\text{nnz}(B)} \sum_{i=1}^{L} \text{nnz}(B_i). \tag{13}$$

In most practical settings, a good value of c ranges between 1 and 4. For details on the choice of $\mathcal{J}_i$, we refer to [2, 4]. The weights, ω_k, are chosen such that the algebraically smoothest vectors have the largest weights in the minimization process. A suitable choice, which we use exclusively in our numerical experiments, is given by

$$\omega_k = \frac{1}{\|Bv\|^2}. \tag{14}$$

In [7], a heuristic argument was proposed suggesting the use of (approximated) singular vectors as test vectors in the nonsymmetric case. The argument follows from the singular value decomposition $B = U\Sigma V^T$. Let the matrix $W = VU^T$ and define the symmetric positive semidefinite matrices $WB = \sqrt{B^T B}$ and $BW = \sqrt{BB^T}$. Then the original nonsymmetric system $Bx = 0$ can be reformulated in two ways as an equivalent symmetric system using WB or BW as the system matrix. The eigenvectors corresponding to the minimal eigenvalues of WB are the *right* singular vectors corresponding to the minimal singular values of B and those for BW are the *left* singular vectors corresponding to minimal singular values of B. Because WB and BW are symmetric positive semi-definite, this fact can be used to derive an approximation property for the original problem involving B, assuming that P and Q are based on singular vectors.

Determining the restriction operator Q for B is equivalent to determining the transpose of the interpolation operator for B^T. This observation naturally leads to a least squares problem similar to (11) for each column of Q, where the weights are now given by $1/\|B^T v\|^2$. This construction, in turn, requires an additional set $\mathcal{U} = \{u^{(k)}, k = 1, \ldots, r\}$, of left test vectors. An efficient multilevel algorithm for constructing the sets $\mathcal{U}$ and $\mathcal{V}$ will now be described.

4.2. Bootstrap construction of test vectors

In this section, a BAMG cycle for computing approximations of singular vectors with small singular values is developed. The algorithm is based on the approach from [4] for approximating small eigenvalues and the corresponding eigenvectors for symmetric operators (where $Q = P^T$, $B_c = P^T BP$). In this approach, a preliminary multigrid hierarchy is constructed and (generalized) eigenvectors

$$P^T BPv = \lambda P^T Pv \tag{15}$$

of the operator on the coarsest grid are determined and interpolated to the finest grid in a suitable way. They are then used as new test vectors for another setup cycle. This approach can be extended to singular vectors by noting that the problem of computing singular vectors is equivalent to solving a symmetric eigenproblem of twice the size, cf. [27]. This has also been observed in [34] in the context of adaptive algebraic multigrid methods for computing singular triplets. In the following, we describe the details of this construction.

To determine suitable approximations to the singular vectors, we consider the equations

$$\left.\begin{array}{rcl} Bv_i &=& \sigma u_i \\ B^T u_i &=& \sigma v_i \end{array}\right\} \quad \text{for } i = 1, \ldots, n, \tag{16}$$

characterizing the singular vectors of B. These equations can be combined into one system of the form

$$\begin{bmatrix} 0 & B \\ B^T & 0 \end{bmatrix} \begin{bmatrix} U & U \\ V & -V \end{bmatrix} = \begin{bmatrix} U & U \\ V & -V \end{bmatrix} \begin{bmatrix} \Sigma & 0 \\ 0 & -\Sigma \end{bmatrix}, \tag{17}$$

where $U = [\, u_1 \mid \cdots \mid u_n \,], V = [\, v_1 \mid \cdots \mid v_n \,]$ and $\Sigma = \operatorname{diag}(\sigma_1, \ldots, \sigma_n)$. Thus the eigenvalues of the symmetric matrix

$$\widehat{B} = \begin{bmatrix} 0 & B \\ B^T & 0 \end{bmatrix}$$

are the singular values of B and the corresponding eigenvectors contain the singular vectors of B. Hence, we can apply the original bootstrap process to the symmetric matrix $\widehat{B}$. To this end define

$$\widehat{P} = \begin{bmatrix} Q^T & 0 \\ 0 & P \end{bmatrix} \tag{18}$$

and observe that in the variational sense an eigenpair $(\widehat{x}, \widehat{\lambda})$ for $\widehat{B}$, satisfying $\widehat{B}\widehat{x} = \widehat{\lambda}\widehat{x}$, corresponds to an eigenpair (x_c, λ_c) of the generalized eigenproblem $\widehat{P^T}\widehat{B}\widehat{P}x_c = \lambda_c\widehat{P^T}\widehat{P}x_c$ on the coarse level, cf. [4, 5]. With the choice (18) of $\widehat{P}$, we obtain the generalized (symmetric) eigenvalue problem on the coarse grid,

$$\begin{bmatrix} 0 & QBP \\ (QBP)^T & 0 \end{bmatrix} \begin{bmatrix} U_c & U_c \\ V_c & -V_c \end{bmatrix} = \begin{bmatrix} QQ^T & 0 \\ 0 & P^TP \end{bmatrix} \begin{bmatrix} U_c & U_c \\ V_c & -V_c \end{bmatrix} \begin{bmatrix} \Sigma_c & 0 \\ 0 & -\Sigma_c \end{bmatrix}, \tag{19}$$

a generalization of (17). Alternatively one can also derive (19) using a Petrov-Galerkin approach as follows: Assume that a singular value σ of B is known and some lower-dimensional subspaces $\mathcal{K}_1, \mathcal{K}_2$ are given. We search for vectors $u_c \in \mathcal{K}_1, v_c \in \mathcal{K}_2$ which approximately satisfy

$$\begin{aligned} Bv_c &= \sigma u_c, \\ B^T u_c &= \sigma v_c. \end{aligned} \tag{20}$$

To identify the approximations u_c, v_c we demand that the residuals $Bv_c - \sigma u_c$ and $B^T u_c - \sigma v_c$ be orthogonal to some subspaces $\mathcal{L}_1, \mathcal{L}_2$ with $\dim \mathcal{L}_i = \dim \mathcal{K}_i$. If we set $\mathcal{K}_1 = \mathcal{L}_2 = \text{range}(Q^T)$, the column span of the matrix Q^T, and $\mathcal{K}_2 = \mathcal{L}_1 = \text{range}(P)$, (20) can be rewritten as

$$\begin{aligned} QBPv_c - \sigma QQ^T u_c &= 0, \\ P^T B^T Q^T u_c - \sigma P^T P v_c &= 0, \end{aligned} \tag{21}$$

which is equivalent to (19). Recursive application of this approach leads to

$$\begin{aligned} B_l v_l - \sigma M_l u_l &= 0, \\ B_l^T u_l - \sigma N_l v_l &= 0, \end{aligned} \tag{22}$$

where $M_l = Q_l \cdots Q_1 Q_1^T \cdots Q_l^T$ and $N_l = P_l^T \cdots P_1^T P_1 \cdots P_l$ are the accumulated interpolation and restriction operators, respectively. On the coarsest level, all possible solutions to (22) are obtained via the generalized eigenvalue problem

$$\begin{bmatrix} 0 & B_L \\ B_L^T & 0 \end{bmatrix} \begin{bmatrix} U_L & U_L \\ V_L & -V_L \end{bmatrix} = \begin{bmatrix} M_L & 0 \\ 0 & N_L \end{bmatrix} \begin{bmatrix} U_L & U_L \\ V_L & -V_L \end{bmatrix} \begin{bmatrix} \Sigma_L & 0 \\ 0 & -\Sigma_L \end{bmatrix}. \tag{23}$$

As long as all P_l and Q_l have full rank, which is always the case in our approach, cf. (12), M_L and N_L are symmetric positive definite and the existence of solutions to (23) is guaranteed, cf. [34].

When the solutions of (23) are computed, they must be transferred back to the finest grid via the interpolation operators, and the quality of approximation is improved by applying the smoothing iteration. To achieve better results, smoothing is not done with respect to homogeneous equations, but with respect to the inhomogeneous equations $B_l v_l = \sigma M_l u_l$ (smoothing for v_l) and $B_l^T u_l = \sigma N_l v_l$ (for u_l) resulting from (22). After a smoothing step for each of u_l and v_l, the subsequent smoothing step uses the smoothed values for v_l and u_l in the right hand sides together with an updated value for the approximate singular value

$$\sigma = \frac{u_l^T B_l v_l}{(u_l^T M_l u_l)^{1/2} (v_l^T N_l v_l)^{1/2}}. \tag{24}$$

This update for σ can be regarded as a generalized Rayleigh quotient. It yields the exact singular value if u_l and v_l are exact generalized singular vectors of B_l in the sense of (22).

To start the above procedure, preliminary interpolation and restriction operators are needed. To get some first information, a few smoothing iterations are applied to the two systems $Bx = 0$ and $B^T x = 0$ to r independent, random initial vectors for each of the matrices, B and B^T. The resulting vectors are then smooth enough to allow for the construction of a first multigrid hierarchy when used as test vectors. This first hierarchy can then be used to produce better approximations to the singular vectors to the r smallest singular vectors of B, which can then in turn be used as new test vectors to construct an even better multigrid hierarchy. In this manner, the method continuously improves itself, which explains why it is referred to as a bootstrap method. We summarize the first cycle of the bootstrap setup in Algorithm 1, where we use the notation $\mathcal{S}_l$ and $\mathcal{S}_l^T$ to denote the smoothing iteration (5) for B_l and B_l^T, respectively, on level l. Subsequent cycles proceed in the same manner except that the first smoothing in lines 5-6 then uses the information which is already available, i.e. it proceeds as in lines 17-18.

Algorithm 1: bamg_mle (BAMG setup, first cycle)

Input: B_l ($B_0 = B$), M_l ($M_0 = I$), N_l ($N_0 = I$), test vectors $\mathcal{U}_l, \mathcal{V}_l$

Output: Triplets $(\Sigma_l, \mathcal{U}_l, \mathcal{V}_l)$ of approximate singular vectors and values of B

if $l = L$ (coarsest grid is reached) **then**

  Determine the r smallest generalized singular triplets $(\Sigma_l, \mathcal{U}_l, \mathcal{V}_l)$ of (23).

else

  for no. of smoothing steps and $k = 1, \ldots, r$

    $\widetilde{v}_l^{(k)} = \mathcal{S}_l(v_l^{(k)}, 0)$. $\widetilde{u}_l^{(k)} = \mathcal{S}_l^T(u_l^{(k)}, 0)$.

    $v_l^{(k)} = \widetilde{v}_l^{(k)}$, $u_l^{(k)} = \widetilde{u}_l^{(k)}$.

  for $m = 1, \ldots, \mu$

    Compute the rows of P_l and columns of Q_l via (11).

    Compute $B_{l+1} = Q_l B_l P_l$.

    Compute $M_{l+1} = Q_l M_l Q_l^T$.

    Compute $N_{l+1} = P_l^T N_l P_l$.

    Set $\mathcal{U}_{l+1} = \{R_l u_l^{(k)} : k = 1, \ldots, r\}$.

    Set $\mathcal{V}_{l+1} = \{R_l v_l^{(k)} : k = 1, \ldots, r\}$.

    $(\Sigma_{l+1}, \mathcal{U}_{l+1}, \mathcal{V}_{l+1}) = \text{bamg_mle}(B_{l+1}, M_{l+1}, N_{l+1}, \mathcal{U}_{l+1}, \mathcal{V}_{l+1}))$.

    Set $(\Sigma_{l+1}, \mathcal{U}_{l+1}, \mathcal{V}_{l+1}) = \{(\sigma_{l+1}^{(k)}, Q_l^T u_{l+1}^{(k)}, P_l v_{l+1}^{(k)}) : k = 1, \ldots, k_v\}$.

    for no. of smoothing steps and $k = 1, \ldots, r$

      $\widetilde{v}_l^{(k)} = \mathcal{S}_l(v_l^{(k)}, \sigma_l^{(k)} M_l u_l^{(k)})$. $\widetilde{u}_l^{(k)} = \mathcal{S}_l^T(u_l^{(k)}, \sigma_l^{(k)} N_l v_l^{(k)})$.

      $v_l^{(k)} = \widetilde{v}_l^{(k)}$, $u_l^{(k)} = \widetilde{u}_l^{(k)}$, $\sigma_l^{(k)} = \frac{(u_l^{(k)})^T B_l v_l^{(k)}}{((u_l^{(k)})^T M_l u_l^{(k)})^{1/2} ((v_l^{(k)} l)^T N_l v_l^{(k)})^{1/2}}$.

5. Application of BAMG to Markov chains

We now go into detail concerning the application of BAMG to computing the steady state vector of a discrete Markov chain. We are interested in solving a linear system with the matrix

$$B = I - A, \tag{25}$$

where A is the transition matrix of a Markov chain and therefore column stochastic, implying that all column sums of B equal zero, i.e., $\mathbf{1}^T B = 0$ with $\mathbf{1} = (1, \ldots, 1)^T$. Since we also assume that A and thus B is irreducible, the left and right nullspaces of B both have dimension one. Nullspaces are identical to the spaces spanned by the singular vectors for the singular value 0. In other words: we know the left singular vector to the smallest singular value, $\sigma_n = 0$, which is $\mathbf{1}$, and we want to compute the unknown corresponding right singular vector $x = v_n$. The bootstrap construction of test vectors of section 4.2 will compute an approximation to v_n together with approximations to further left and right singular vectors belonging to small singular values.

It is reasonable to try to take advantage of the additional information that $\mathbf{1}$ is a left singular vector for the singular value 0. We do so by preserving the property $\mathbf{1}^T B = 0$ on the coarser levels. On all levels we then know the left singular vector for the smallest singular value exactly. To achieve this, we use the exact singular vector $\mathbf{1}$ as a test vector within the least squares process (11) to define Q. For $u^{(1)} = \mathbf{1}$ this results in the weight $\omega_1 = \infty$ with respect to B^T according to (14),

$$\omega_1 = \frac{1}{\|B^T \mathbf{1}\|^2} = \frac{1}{0} = \infty. \tag{26}$$

Therefore, the least squares functional $\mathcal{L}$ will attain the value ∞ as long as the vector $\mathbf{1}$ is not interpolated exactly. As a modification to the bootstrap process for Markov chains we thus require that $\mathbf{1}$ be interpolated exactly and set $0 \cdot \infty = 0$, so that (11) turns into a least squares problem for the other test vectors $u^{(2)}, \ldots, u^{(r)}$, subject to the restriction $Q_i^T \mathbf{1} = \mathbf{1}$ for all columns i of Q. Restricted least squares problems of this kind can easily be solved by known techniques or transformed into standard least squares problems, see [27]. Overall, the restrictions for the individual columns of Q

imply that the computed restriction operator will satisfy $\mathbf{1}^T Q = \mathbf{1}^T$, leading to $\mathbf{1}^T B_c = \mathbf{1}^T QBP = \mathbf{1}^T BP = 0$. In our numerical experiments we noticed that preserving the constant in this manner improves the quality of the coarse grid operators when compared to the simple averaging (the non-zero structure of the restriction operator Q is chosen as the transpose of the non-zero structure of P with all entries in one column set to $\frac{1}{c}$, where c is the number of non-zero entries in this column) from [2, 12, 3, 21], another attempt to preserve structural properties for the coarse grid operators.

6. BAMG-preconditioned GMRES

Each setup cycle, as described in section 4, involves recomputing the operators P, Q, and then $B_c = QAP$, and is therefore computationally quite costly. As demonstrated in [2] and then also later explored in [34], the use of a second, *additive* phase—where one iterates with a fixed multigrid hierarchy to improve the quality of the approximation to the state vector obtained so far—can result in substantial speed up. This idea is motivated by the fact that the multigrid hierarchy developed by the setup cycles approximates a rich subspace spanned by singular vectors of B with small singular values (in addition to the one with singular value 0, the state vector) which makes the resulting hierarchy well suited for forming an effective preconditioner for GMRES. Specifically, let $E_0 = I - \widetilde{C}B$ denote the two-grid error propagation operator (9). Then, given the approximation $x^{(0)}$ to the steady state vector x, computed in the setup phase, and an arbitrary initial guess $e^{(0)}$ for the residual equation

$$Be = -Bx^{(0)}, \tag{27}$$

one application of the multigrid V-cycle preconditioner yields a new iterate

$$e^{(1)} = (I - \widetilde{C}B)e^{(0)} - \widetilde{C}Bx^{(0)}. \tag{28}$$

Thus, if we start with initial guess $e^{(0)} = 0$, one multigrid iteration corresponds to a multiplication with the implicitly defined preconditioner $C = \widetilde{C}$, allowing us to use GMRES on the preconditioned residual equation

$$CBx = -CBx^{(0)}. \tag{29}$$

For singular linear systems, convergence of the GMRES iteration can in general not be guaranteed, even when the linear system is consistent. The following theorem on the convergence of GMRES for singular systems summarizes results from [35]. Recall that B is said to have index 1 if $\text{range}(B) \cap \text{null}(B) = 0$.

Theorem 1. *Assume that the matrix B is singular and has index 1. Then the GMRES iterates for the singular system $Bx = b$ with $b \in \text{range}(B)$ converge towards the solution $x = x^{(0)} + B^D(b - Bx^{(0)})$, where $x^{(0)}$ is the starting vector and B^D denotes the Drazin inverse of B.*

For the matrix $B = I - A$ with A the transition matrix of an irreducible Markov chain, we have $\text{null}(B) = \langle x \rangle$ and $\text{range}(B) = \langle \mathbf{1} \rangle^\perp$ where x is the steady state vector of the Markov chain. Because all components of x are strictly positive, x can not be orthogonal to the constant vector $\mathbf{1}$ and therefore B has index 1. Since the residual equation (27) is consistent, Theorem 1 gurantees the convergence of the GMRES iteration.

It is not immediately clear that the preconditioned matrix CB also has index 1, since CB in general will not have column sums all equal to zero, so that we will no longer have that $\text{range}(CB) = \langle \mathbf{1} \rangle^\perp$.

However in all our numerical experiments from section 7 we never experienced a breakdown of GMRES, indicating that it is highly unlikely in practice that CB has index > 1. We proceed by formulating results which, to a certain extent, quantify the acceleration of convergence due to the preconditioning. Theorem 2 below recalls a result from [36] and [37] on the convergence of GMRES in the non-singular case. It uses the field of values $\mathcal{F}(B)$ of B defined as

$$\mathcal{F}(B) = \{\langle Bx, x \rangle : \langle x, x \rangle = 1, x \in \mathbb{C}^n\}.$$

$\mathcal{F}(B)$ is a compact set containing all the eigenvalues of B, see [38]. We denote by $\nu(\mathcal{F}(B))$ the distance of the field of values to the origin, i.e.

$$\nu(\mathcal{F}(B)) = \min_{z \in \mathcal{F}(B)} \|z\|.$$

Theorem 2. *The residuals $r_k = b - Bx_k$ of the GMRES iterates x_k satisfy*

$$\frac{\|r_k\|}{\|r_0\|} \leq \left(1 - \nu(\mathcal{F}(B))\nu(\mathcal{F}(B^{-1}))\right)^{k/2}. \tag{30}$$

It is easy to see that $\nu(\mathcal{F}(B)) > 0$ if and only if $\nu(\mathcal{F}(B^{-1})) > 0$, and that $\nu(\mathcal{F}(B)) \cdot \nu(\mathcal{F}(B^{-1})) < 1$. Thus the theorem predicts progress in the GMRES iteration in every step provided $\nu(\mathcal{F}(B)) > 0$, and the larger $\nu(\mathcal{F}(B)) \cdot \nu(\mathcal{F}(B^{-1}))$ is, the faster the predicted progress is. It has been shown in [36, 39] that (30) is an improvement over the estimate from [40] (see also [41, 42])

$$\frac{\|r_k\|}{\|r_0\|} \leq \left(1 - (\nu(\mathcal{F}(B))/\|B\|)^2\right)^{k/2} \tag{31}$$

in the sense that $1 - \nu(\mathcal{F}(B))\nu(\mathcal{F}(B^{-1}) \leq 1 - (\nu(\mathcal{F}(B))/\|B\|)^2$.

In order to, at least partly, explain the fast convergence of the BAMG-preconditioned GMRES iteration we now relate the GMRES iteration for the singular system to one for a non-singular system of smaller dimension via the following theorem based on [43, Thm.2.9].

Theorem 3. *Assuming the singular matrix B has index 1 and* $\dim(\text{range}(B)) = m < n$. *Let $\Pi \in \mathbb{R}^{n\times m}$ have orthonormal columns which span* range(B) *and define*

$$\widehat{B} = \Pi^T B \Pi \in \mathbb{R}^{m\times m}.$$

Assume that $b \in$ range(B) and consider the GMRES iterates x_k for the system $Bx = b$ with starting vector x_0. Decompose $x_0 = x_0^1 + x_0^2$ with $x_0^1 \in$ range(B), $x_0^2 \in$ null(B) and consider the GMRES iterates for the system $\widehat{B}\widehat{x} = \widehat{b}$ with $\widehat{b} = \Pi^T b$ and $\widehat{x_0} = \Pi^T x_0^1$. Then we have

$$x_k = x_0^2 + \Pi\widehat{x_k} \text{ and } r_k = b - Bx_k = \Pi(\widehat{B}(\widehat{b} - \widehat{x_k}) = \Pi\widehat{r_k}. \tag{32}$$

In particular, if $\nu(\mathcal{F}(\widehat{B})) > 0$ we have

$$\frac{\|r_k\|}{\|r_0\|} \leq \left(1 - \nu(\mathcal{F}(\widehat{B}))\nu(\mathcal{F}(\widehat{B}^{-1}))\right)^{k/2}.$$

Proof. The equality (32) was established in [43]. The subsequent estimate follows from Theorem 2.

Using this theorem we can substantiate the experimentally observed fast convergence of the BAMG preconditioned GMRES iteration as compared to plain GMRES by comparing the fields of values of the respective projected matrices $\widehat{CB}$ and $\widehat{B}$, respectively. Note that since $\langle\widehat{B}y, y\rangle = \langle B\Pi y, \Pi y\rangle$ and $\|\Pi y\| = \|y\|$ we have

$$\mathcal{F}(\widehat{B}) = \{\langle\widehat{B}y, y\rangle : \|y\| = 1, y \in \mathbb{C}^m\} = \{\langle Bz, z\rangle : \|z\| = 1, z \in \mathbb{C}^n, z = \Pi y\} \subseteq \mathcal{F}(B).$$

Figure 2 shows this comparison for three of the (smaller) examples of section 7. The first column depicts the field of values of the original matrix B, together with its eigenvalues. The second and third column give the same information for the projected matrix $\widehat{B}$ and the projected preconditioned matrix $\widehat{CB}$. The figure shows that the field of values almost remains unchanged when we go from B to $\widehat{B}$. We note that 0 must lie in the topological interior of $\mathcal{F}(B)$, since if it were contained in the boundary, we would have null(B) = null(B^T) by a result of [44], see also [38], i.e. $x = \mathbf{1}$. This would mean that A is doubly stochastic which is not the case in our examples. An additional information, which due to the scaling cannot be seen in the figure, is that the field of values $\mathcal{F}(\widehat{B})$ still contains 0 in its interior in all three examples, so that Theorem 2 cannot be used to predict a bound for the convergence speed of (non-preconditioned) GMRES. For the preconditioned systems, $\mathcal{F}(\widehat{CB})$ is nicely bounded away from 0, and thus Theorem 3 predicts a rapidly convergent method which is what we observe in our numerical experiments.

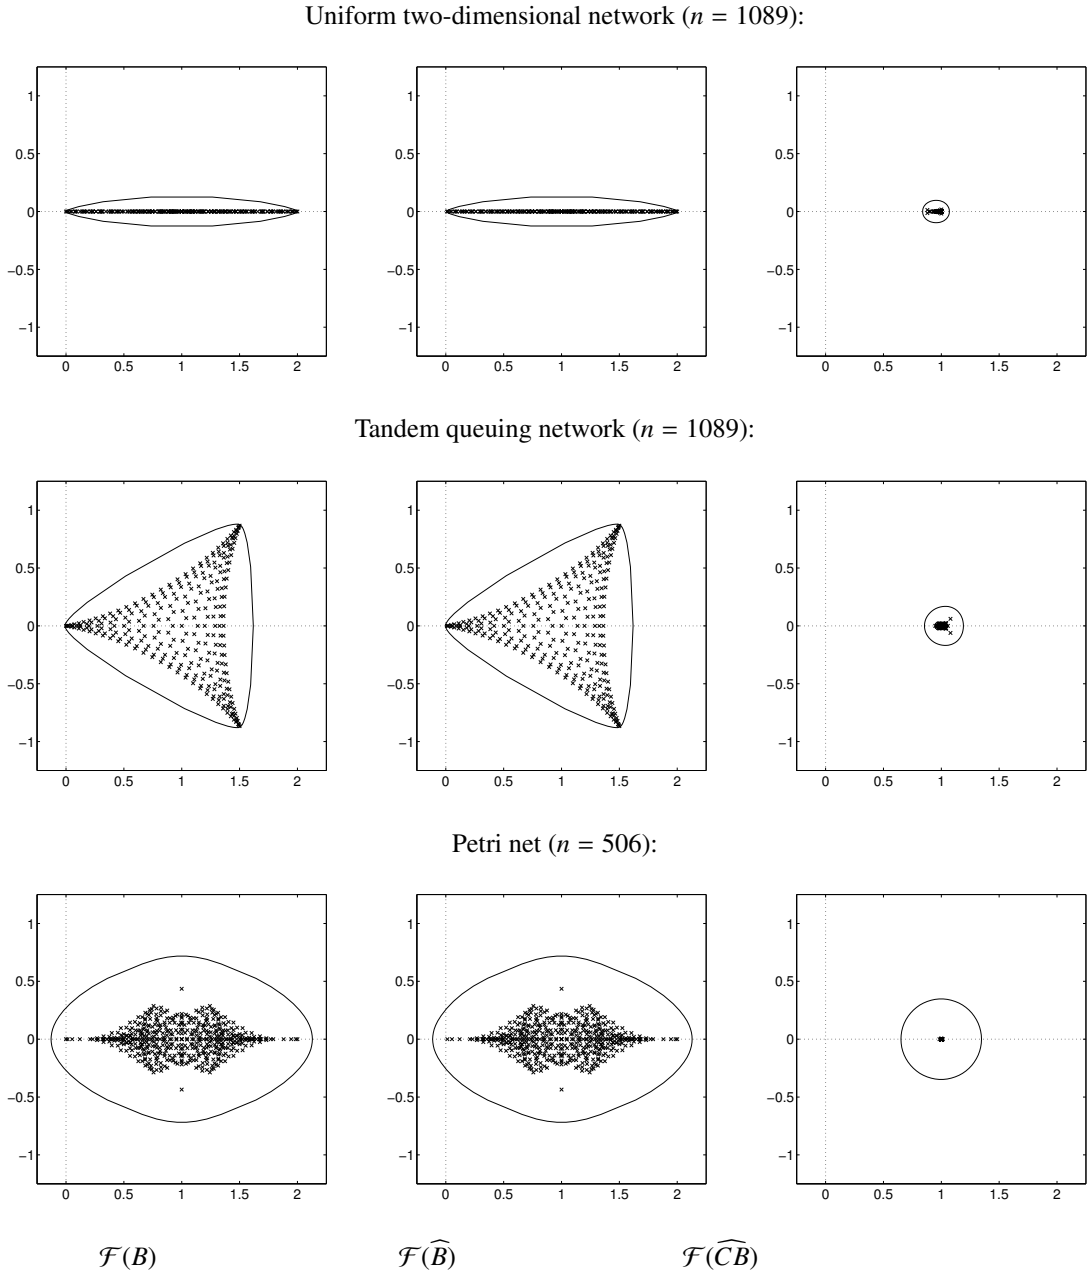

Figure 2: Spectrum and field of values of the matrices from three different test problems from section 7 before and after preconditioning by one multigrid V-cycle.

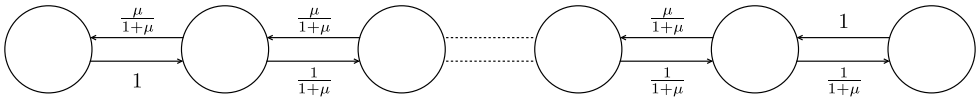

Figure 3: One-dimensional birth-death chain

n	1025	2049	4097	8193	16385	32769	65537
GMRES(50)	> 1000	> 1000	> 1000	> 1000	> 1000	> 1000	> 1000
BAMG + GMRES	3	3	3	4	4	4	5
BAMG^2 + GMRES	1	1	1	1	1	2	2

Table 1: Iteration numbers for the one-dimensional birth-death chain

7. Numerical examples

In this section, we report the results of numerical tests performed with the proposed method for a variety of Markov chains. The test problems are chosen such that many different scenarios occur: We consider (1) problems with a regular geometry as well as problems without any exploitable structure, (2) problems in which all transition probabilities are within the same order of magnitude as well as problems with highly varying probabilities, and (3) structurally symmetric problems as well as nonsymmetric ones. For each test example, we report the number of iterations needed to reduce the scaled residual $\|Bx^{(k)}\|/\|x^{(k)}\|$ for the iterate $x^{(k)}$ to 10^{-7}, both for the BAMG preconditioned GMRES method and the GMRES method without preconditioning. For the latter, to limit memory and computational cost, we restart after every 50 iterations, i.e. we perform GMRES(50). As a complement to the operator complexity o_c from (13), we also report the grid complexity $o_g = \frac{1}{n}\sum_{i=1}^{L} n_i$. These measures are commonly used in the multigrid literature to give an indication of how efficiently one cycle of the multigrid method can be executed. The lower the complexities, the more efficient a single cycle will be. The standard parameters for the method are three pre- and post-smoothing iterations, which is always the weighted Jacobi iteration with $\omega = 0.7$, and $r = 8$ test vectors, unless explicitly stated otherwise. Most of the test cases are taken either from [2] or from [3].

The first test example, a birth-death chain, has a regular one-dimensional structure. From each node, only transitions to its left and right neighbor are possible as depicted in Figure 3, which leads to a tridiagonal matrix B. We choose the parameter $\mu = 0.96$ as in [2, 3].

For this example geometric coarsening (i.e., choosing every other variable as a coarse variable) is performed (until a problem size $n_c < 30$ is reached) and interpolation is done only from direct neighbors (with at most $c = 2$ interpolatory points), resulting in grid and operator complexities that are slightly below 2. The required iteration numbers for this problem are shown in Table 1. In the first row, the required number of iterations of non-preconditioned GMRES is reported. The second and third rows contain results obtained by performing one or two BAMG setup cycles, respectively, and then using the fixed hierarchy as a preconditionener for GMRES. The given integer value indicates the number of GMRES iterations needed in addition to the BAMG setup cycles to achieve the prescribed stopping tolerance, i.e., 1 means that one or two setup cycles and one iteration of preconditioned GMRES we performed. The iteration numbers clearly show that preconditioning improves the convergence behavior of GMRES substantially and that the method scales for the given range of problem sizes. The use of two setup cycles gives another noticeable improvement over just one cycle. This is to be expected since for two setup cycles the transfer operators are constructed with respect to the approximate singular vectors from the previous cycle, as opposed to just smoothed random vectors when only one cycle is used.

Due to the different transition probabilities to the left and to the right, the solution vector x contains entries with highly varying order of magnitude, ranging from $O(\mu^n)$ on the left end of the chain to $O(1)$ on the right end. This is one reason why this test example is difficult to solve for some methods, as can be seen from the results for the non-preconditioned GMRES method in Table 1. For the method considered in [2] that, while the residual can be reduced below the prescribed tolerance in a few iterations, the computed solution contains small negative entries, which should not be the case. In the preconditioned method proposed here, we do not observe this behavior.

In the next two test examples, the Markov chains have a regular two-dimensional geometric structure. The first two-dimensional problem we consider is a uniform network, as depicted in Figure 4a. All transition probabilities from

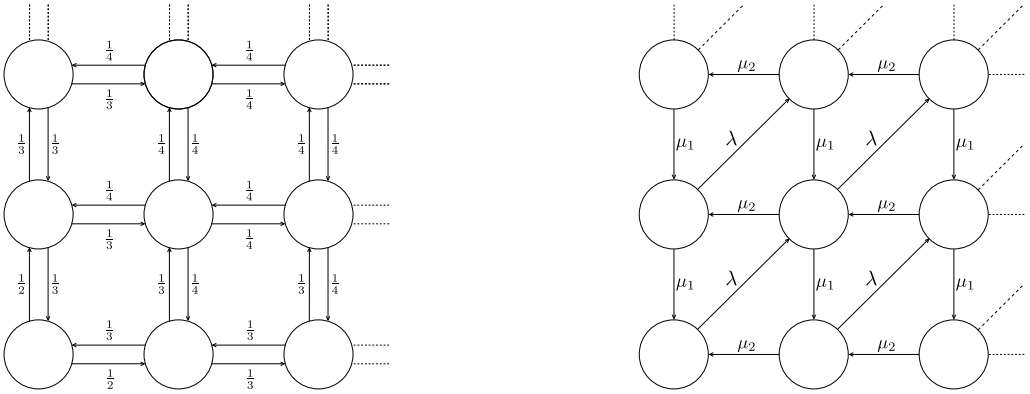

Figure 4: Uniform two-dimensional network (left) and two-dimensional tandem queuing network

n	1089	4225	9409	16641	25921	66049
GMRES(50)	52	89	105	170	200	240
BAMG + GMRES	6	6	6	7	8	10
BAMG^2 + GMRES	4	4	4	5	5	6

Table 2: Iteration numbers for the uniform two-dimensional network

one state are equal to the reciprocal of the number of neighboring states. We again use our standard parameters and geometric coarsening. In the two-dimensional case this means that every other node in both directions is chosen as a coarse level node. We allow interpolation from up to $c = 4$ neighboring coarse level variables which leads to grid and operator complexities that are bounded by 1.4 and 1.8, respectively. Coarsening is done down to a problem size of 17^2. The results obtained for this problem are reported in Table 2.

As in the one-dimensional test case, we see that BAMG preconditioning needs far less iterations than non-preconditioned GMRES. The scaling behaviour of the BAMG preconditioned method is again very favorable, although we now need more than one iteration also in case of two setup cycles.

The second test case with a regular two-dimensional geometric structure is the tandem queuing network example from [45] which is depicted in Figure 4b. The Markov chain considered here is the embedded discrete Markov chain of a continuous-time homogeneous Markov chain that describes a queuing system where each customer has to be served at two stations with one server each. In the continuous time case, the service time at station 1 and 2 is exponentially

n	1089	4225	9409	16641	25921	66049
GMRES(50)	165	295	410	570	750	1030
BAMG + GMRES	6	6	8	9	10	12
BAMG^2 + GMRES	4	4	4	5	5	6

Table 3: Iteration numbers for the tandem queuing network

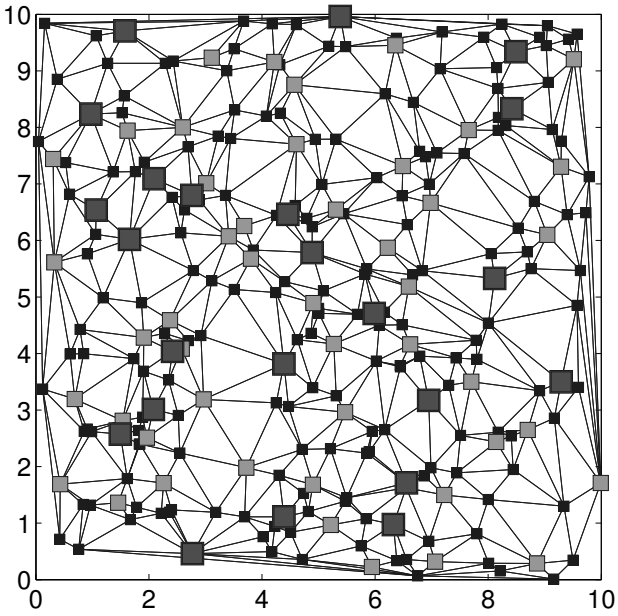

Figure 5: Random planar graph with 256 nodes and coarsening by compatible relaxation. The smallest nodes are the initial grid, the medium size nodes are selected for the second level and the largest nodes for the third level

n	1024	4096	8192	16384	32768	65536
GMRES(50)	50	80	105	115	135	170
BAMG + GMRES	8 (2)	10 (2)	12 (3)	15 (3)	19 (4)	22 (4)
BAMG^2 + GMRES	3 (2)	3 (2)	4 (3)	5 (3)	5 (4)	6 (4)

Table 4: Iteration numbers and number of levels for the random planar graphs

distributed with mean value $\frac{1}{\mu_1}$ and $\frac{1}{\mu_2}$, respectively, and customers arrive according to a Poisson process with rate λ. In the discrete case, this leads to transition probabilities μ_1, μ_2 and λ, as in Figure 4b. The diagonal transitions describe the arrival of a new customer, the horizontal and vertical arrows describe that a customer has been served at station 1 or 2. We chose the parameters to be $\lambda = \frac{11}{31}, \mu_1 = \frac{10}{31}$ and $\mu_2 = \frac{10}{31}$, as in [2, 3]. In contrast to the test cases so far, the matrix of the tandem queuing network is structurally nonsymmetric and has a complex spectrum. We test the method with the same parameters as we used for the uniform two-dimensional network and the results are given in Table 3. The behavior is almost the same as for the uniform two-dimensional network, although the problem is much harder to solve for the non-preconditioned GMRES iteration. This again underlines the high quality of the adaptively computed multigrid hierarchy.

Our last two test examples do not exhibit any regular geometric structure, so that the coarsening has to be done algebraically. In the tests we use a compatible relaxation algorithm [33]. We first consider random walks on planar graphs. These planar graphs are generated, as suggested in [21], by choosing n random points in the unit square and then constructing a Delaunay triangulation of these points. Transitions between states are then restricted to the edges of the triangles, with transition probabilities again chosen as the reciprocals of the number of outgoing edges for each node. For these problems, we use $r = 8$ test vectors, coarsen down to a grid size $n_c < 500$ and allow interpolation from up to $c = 3$ neighboring nodes. The number of pre- and postsmoothing iterations are both increased to 5. The results of the method are given in Table 4. The numbers in brackets report the number of levels used to reach a problem size of less than 100 variables on the coarsest level. This results in grid and operator complexities which are bounded

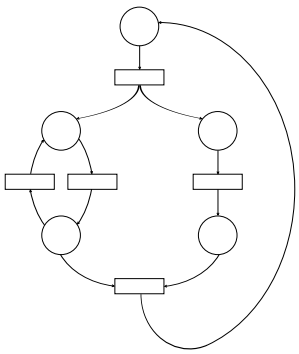

Figure 6: Petri net with five places and five transitions

n	1496	3311	10416	23821
GMRES(50)	72	90	95	180
BAMG + GMRES	6	6	6	6
BAMG^2 + GMRES	5	5	5	6

Table 5: Iteration numbers for the Petri net

by 1.3 and 1.5, respectively, for all considered problem sizes. The results show a significant improvement compared to the eigenvector based BAMG approach in [2], where a substantial increase in iteration numbers was observed for increasing problem sizes. With the proposed singular vector based BAMG approach we observe now only a mild dependence on the number of iterations on the problem size, especially when using two setup cycles.

The last test example is a Petri net taken from [46]. It was, e.g., also considered in [24] and [19]. Petri nets are used for the description of distributed systems and consist of places and transitions. In the simple Petri net depicted in Figure 6, the circles represent places and the bars represent transitions. The places can be seen as conditions under which certain transitions can occur: A transition with ingoing edges from two places can only occur when both places are marked by at least one *token*. When the transitions occur, the tokens will disappear and in each place reached by an outgoing edge of the transition, a new token will appear. This way, starting from an initial marking, the tokens will travel through the different places of the Petri net. A Petri net can be transformed into a discrete Markov chain as follows (cf. [47]): Each marking of the Petri net that is reachable from a given initial marking corresponds to one state of the Markov chain. Transitions occur between all markings that can be converted into each other by the firing of one transition. The transition probabilities can then be calculated from the known firing probabilities of the transitions of the Petri net. The size of the resulting Markov chain depends on the number of tokens in the initial marking and grows very fast. When the initial marking consists of only one token in the uppermost node in the Petri net from Figure 6, the Markov chain will have only 5 states, for 10 tokens in the same node the Markov chain has 506 states, and for 30 tokens it already consists of 10416 states, although the underlying Petri net has only 5 places and 5 transitions. The transition matrix of the resulting Markov chain is structurally nonsymmetric with complex spectrum, as depicted in Figure 2. We chose the firing probabilities of the transitions to be the same as in [24] and [19]. Since there is no underlying geometric structure, coarsening is again done by compatible relaxation. We use $r = 10$ test vectors and the other parameters of the algorithm are the same as for the planar graph problems. The results of the tests are reported in Table 5 and the grid and operator complexities are bounded by 1.7 and 2.5, respectively.. The problem sizes correspond to a Petri net with an initial marking of 15, 20, 30 and 40 tokens in the uppermost node, respectively. While the two versions of the preconditioned GMRES iteration scale for the problem sizes considered, we see no significant reduction in the iteration numbers when using two setup cycles instead of one.

8. Conclusions

In this paper, we developed a bootstrap algebraic multigrid framework for non-symmetric matrices. We demonstrated that using singular vectors instead of eigenvectors in the adaptive construction of transfer operators leads to

a new, more robust approach for the targeted Markov problems. The approach is shown to be especially suitable in the context of Markov chains because natural assumptions on the coarse grid matrix, e.g., singularity and column sum zero, can automatically be fulfilled when using the non-symmetric bootstrap setup. We showed that this can be accomplished by modifying the least squares interpolation to allow for linear constraints.

We experimentally explained the fast convergence of the preconditioned system by relating the GMRES iteration for the singular system to the one for an equivalent, projected non-singular system with a system matrix that has its field of values well separated from 0. Standard results on the speed of convergence of GMRES in the non-singular case can then be used to explain the fast convergence of the method.

The computational examples show the superiority of the singular vector based method over the eigenvector based method from [2]. The method further reduces the iteration count and yields near-optimal scaling behavior for the considered test problems, especially when two setup cycles are used, therefore taking full advantage of the improved test vectors. Thus we expect the new method to be particularly efficient for large problem sizes. When comparing computational cost in our numerical examples one should be aware that one iteration with the BAMG preconditioner is roughly equivalent to $(s_{pre} + s_{post})o_c$ matrix-vector multiplications, where s_{pre} and s_{post} is the number of pre- and post-smoothing iterations, respectively, and o_c is the operator complexity. For each BAMG setup cycle, this cost has to be multiplied by $2r$, r being the number of test vectors in $\mathcal{U}$ and $\mathcal{V}$. Measured in this manner, our numerical examples indicate that for the still quite moderate problem sizes considered there we already have a substantial gain of the new BAMG preconditioned method over non-preconditioned GMRES for the more difficult problems, i.e. the death-birth chain, the tandem queueing network and the Petri net.

References

[1] A. Berman, R. J. Plemmons, Nonnegative Matrices in the Mathematical Sciences, SIAM, 1994.

[2] M. Bolten, A. Brandt, J. Brannick, A. Frommer, K. Kahl, I. Livshits, A bootstrap algebraic multilevel method for Markov chains, SIAM J. Sci. Comput. 33 (6) (2011) 3425–3446.

[3] H. de Sterck, T. A. Manteuffel, S. F. McCormick, J. Ruge, K. Miller, J. Pearson, G. Sanders, Smoothed aggregation multigrid for Markov chains, SIAM J. Sci. Comput. 32 (1) (2010) 40–61.

[4] A. Brandt, J. Brannick, K. Kahl, I. Livshits, Bootstrap AMG, SIAM J. Sci. Comput. 33 (2) (2011) 623–632.

[5] K. Kahl, Adaptive algebraic multigrid for lattice QCD computations, Ph.D. thesis, Bergische Universität Wuppertal, Fachbereich Mathematik und Naturwissenschaften, Wuppertal (2009).

[6] T. A. Manteuffel, S. F. McCormick, M. Park, J. Ruge, Operator-based interpolation for bootstrap algebraic multigrid, Numer. Linear Algebra Appl. 17 (2–3) (2010) 519–537.

[7] M. Brezina, T. A. Manteuffel, S. F. McCormick, J. Ruge, G. Sanders, Towards adaptive smoothed aggregation (αSA) for nonsymmetric problems, SIAM J. Sci. Comput. 32 (1) (2010) 14–39.

[8] J. E. Moyal, Stochastic processes and statistical physics, J. Roy. Statist. Soc., Ser. B. 11 (1949) 150–210.

[9] M. F. Neuts, Markov chains with applications in queueing theory, which have a matrix-geometric invariant probability vector, Adv. in Appl. Probab. 10 (1) (1978) 185–212.

[10] U. R. Krieger, B. Muller-Clostermann, M. Sczittnick, Modeling and analysis of communication systems based on computational methods for Markov chains, IEEE J. Sel. Areas Commun. 8 (9) (1990) 1630–1648.

[11] G. Benoît, Application of Markov chains in an interactive information retrieval system, Inf. Process. Manage. 41 (4) (2005) 843–857.

[12] H. de Sterck, T. A. Manteuffel, S. F. McCormick, Q. Nguyen, J. Ruge, Multilevel adaptive aggregation for Markov chains, with application to web ranking, SIAM J. Sci. Comput. 30 (5) (2008) 2235–2262.

[13] A. N. Langville, C. D. Meyer, Google's PageRank and Beyond: The Science of Search Engine Rankings, Princeton University Press, 2006.

[14] M. Haviv, Aggregation/disaggregation methods for computing the stationary distribution of a Markov chain, SIAM J. Numer. Anal. 24 (4) (1987) 952–966.

[15] J. R. Koury, D. F. McAllister, W. J. Stewart, Iterative methods for computing stationary distributions of nearly completely decomposable Markov chains, SIAM J. Algebraic Discrete Methods 5 (2) (1984) 164–186.

[16] U. R. Krieger, On a two-level multigrid solution method for finite Markov chains, Linear Algebra Appl. 223–224 (0) (1995) 415–438.

[17] H. A. Simon, A. Ando, Aggregation of variables in dynamic systems, Econometrica 29 (1961) 111–138.

[18] Y. Takahashi, A lumping method for numerical calculations of stationary distributions of Markov chains, Tech. rep., B-18, Department of Information Sciences, Tokyo Institute of Technology (1975).

[19] G. Horton, S. T. Leutenegger, A multi-level solution algorithm for steady-state Markov chains, in: B. D. Gaither (Ed.), Proceedings of the ACM SIGMETRICS 1994 Conference on Measurement and Modeling of Computer Systems, 1994, pp. 191–200.

[20] G. Horton, S. T. Leutenegger, On the utility of the multi-level algorithm for the solution of nearly completely decomposable Markov chains, Tech. rep., Institute for Computer Applications in Science and Engineering (ICASE) (1994).

[21] H. de Sterck, K. Miller, G. Sanders, M. Winlaw, Recursively accelerated multilevel aggregation for Markov chains, SIAM J. Sci. Comput. 32 (2010) 1652–1671.

[22] E. Virnik, An algebraic multigrid preconditioner for a class of singular M-matrices, SIAM J. Sci. Comput. 29 (5) (2007) 1982–1991.

[23] H. de Sterck, T. A. Manteuffel, K. Miller, G. Sanders, Top-level acceleration of adaptive algebraic multilevel methods for steadystate solution to Markov chains, Adv. Comput. Math. 35 (2010) 375–403.

[24] H. de Sterck, K. Miller, G. Sanders, Iterant recombination with one-norm minimization for multilevel Markov chain algorithms via the ellipsoid method, Comput. Vis. Sci. 14 (2011) 51–65.
[25] W. Hackbusch, Iterative Solution of Large Sparse Systems of Equations, Springer-Verlag, New York, NY, USA, 1994.
[26] S. Kaczmarz, Angenäherte Auflösung von Systemen linearer Gleichungen, Bull. Internat. Acad. Polon. Sci. Lettres A (1937) 355–357.
[27] G. H. Golub, C. F. V. Loan, Matrix Computations, Johns Hopkins University Press, 1996.
[28] O. Axelsson, P. Vassilevski, Algebraic multilevel preconditioning methods i, Numer. Math. 56 (1989) 157–177.
[29] U. Trottenberg, C. Osterlee, A. Schüller, Multigrid, Academic Press, 2001.
[30] J. Ruge, K. Stüben, Algebraic multigrid, Multigrid Methods (McCormick, S.F., ed.).
[31] A. Brandt, General highly accurate algebraic coarsening, Electron. Trans. Numer. Anal. 10 (2000) 1–20.
[32] J. Brannick, Adaptive algebraic multigrid coarsening strategies, Ph.D. thesis, University of Colorado at Boulder (2005).
[33] J. Brannick, R. Falgout, Compatible relaxation and coarsening in algebraic multigrid, SIAM J. Sci. Comput. 32 (3) (2010) 1393–1416.
[34] H. de Sterck, A self-learning algebraic multigrid method for extremal singular triplets and eigenpairs, SIAM J. Sci. Comput. 34 (4) (2012) A2092–A2117.
[35] I. C. F. Ipsen, C. D. Meyer, The idea behind Krylov methods, Amer. Math. Monthly 105 (1998) 889–899.
[36] M. Eiermann, O. G. Ernst, Geometric aspects of the theory of Krylov subspace methods, Acta Numerica 10 (2001) 251–312.
[37] G. Starke, Field-of-values analysis of preconditioned iterative methods for nonsymmetric elliptic problems, Numer. Math. 78 (1997) 103–117.
[38] R. A. Horn, C. R. Johnson, Topics in Matrix Analysis, Cambridge University Press, Cambridge, 1994.
[39] J. Liesen, P. Tichý, The field of values bound on ideal gmresArXiv:1211.5969v2 [math.NA].
[40] H. C. Elman, Iterative methods for large sparse nonsymmetric linear systems of equations, Ph.D. thesis, Department of Computer Science, Yale University (1982).
[41] Y. Saad, Iterative Methods for Sparse Linear Systems, 2nd Edition, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2003.
[42] Y. Saad, M. H. Schultz, GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems, SIAM J. Sci. Comput. 7 (3) (1986) 856–869.
[43] K. Hayami, M. Sugihara, A geometric view of Krylov subspace methods on singular systems, Numer. Linear Algebra Appl. 18 (3) (2011) 449–469.
[44] R. Kippenhahn, Über den Wertevorrat einer Matrix, Mathematische Nachrichten 6 (3-4) (1951) 193–228. doi:10.1002/mana.19510060306.
URL http://dx.doi.org/10.1002/mana.19510060306
[45] W. J. Stewart, Introduction to the numerical solution of Markov chains, Princeton University Press, 1994.
[46] M. K. Molloy, Performance analysis using stochastic Petri nets, IEEE Trans. Comput. 31 (9) (1982) 913–917.
[47] H. Motameni, A. Movaghar, M. Siasifar, H. Montazeri, A. Rezaei, Analytic evaluation on Petri net by using Markov chain theory to achieve optimized models, World Applied Sciences Journal 3 (2008) 504–513.